



(19) Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) Publication number:

0 439 646 A1

(12)

EUROPEAN PATENT APPLICATION

(21) Application number: 90101759.0

(51) Int. Cl. 5: H04L 12/44

(22) Date of filing: 30.01.90

The title of the invention has been amended
(Guidelines for Examination in the EPO, A-III,
7.3).

(43) Date of publication of application:
07.08.91 Bulletin 91/32

(84) Designated Contracting States:
DE FR GB IT

(71) Applicant: Hewlett-Packard Company
Mail Stop 20 B-O, 3000 Hanover Street
Palo Alto, California 94304(US)

(72) Inventor: McKeown, Nicholas William, c/o
Hewlett-Packard Ltd
Network & Comm. Lab. Inf. Syst. Centr.,
Filton Rd.
Stoke Gifford, Bristol BS12 6QZ(GB)


(74) Representative: Squibbs, Robert Francis et al
Hewlett-Packard Limited Cain Road
Bracknell, Berkshire RG12 1HN(GB)

(54) Optical star network protocol and system with minimised delay between consecutive packets.

(57) A system to allow a network of stations (100) such as computers arranged in a star formation, to communicate with each other. The system has a protocol whereby stations (100) are granted access to the network, the protocol being tailored to star networks. Each station (100) may send a packet of data in each cycle of operation, provided that it indicated in a corresponding reservation slot (112), that it wishes to transmit. By examining the reservation slots (112) associated with stations (100) earlier in the cycle, each station (100) can determine the station (100) which is immediately preceding the one station amongst the stations which are to transmit data (i.e. have indicated in the corresponding reservation slot). On the basis of that determination the one station (100) transmits its data packet promptly after that preceding station (100) has transmitted its data packet; at least in certain cases, the station initiates transmission prior to having finished receiving the packet transmitted by the preceding station but with a timing to avoid data collision. Arrangements are shown whereby a data receiving station (100) can acknowledge receipt in a suitable acknowledgement slot (114) and wherein a busy station (100) can automatically reserve space. A variation to the protocol is given wherein reservations are note necessary, thereby speeding up operation in certain

operating conditions.

FIG 5



EP 0 439 646 A1

COMMUNICATION SYSTEM

The present invention relates to a communication system, and in particular to a digital communication system between a plurality of stations such as digital computers. The present invention is particularly, but not exclusively, applicable to such communication systems in the form of a star in which peripheral stations are connected directly or indirectly to a central node.

Any communications system must utilize a protocol for determining the sequence and nature of the permitted communication over the system by any station on it. In the past, protocols for star networks have generally been those which were developed for other networks such as rings or buses. Thus, it is known to operate star networks on the basis of token passing protocols that were originally developed for ring networks. With such protocols, a logical "token" is passed from station to station and any station "holding" the token is allowed to transmit data. A disadvantage of token-passing protocols is that the gap (measured in bits) between the end of transmission by one station, and the start of transmission by the next station to receive the token, will increase with the data rate. This overhead caused by the passing of the token, places both a lower bound on the access delay and an upper bound on the efficiency of utilisation of the system, and in addition limits the efficiency under low loading. Thus, token passing protocols are not efficient on a star network. Furthermore, in the case of lost or duplicate tokens, a considerable distributed management problem is created.

Star networks have also been operated using random access protocols, such as CSMA/CD (Carrier Sense Multiple Access/Collision Detect). However, with such protocols, the minimum length of the transmission by a station must increase with the data rate, and this produces inefficiency. An increase in data rate is not compensated for by a reduction in access delay, since the number of "collisions" between data from different stations will remain approximately constant.

Another known type of protocol for broadcast communications systems is the so-called bit-map protocol; such protocols have previously been used, for example, in packet radio systems. Typically, in a bit-map protocol the stations reserve the use of subsequent slots in an operating cycle by setting an appropriate bit in a reservation field (or bit-map); the stations thereafter use their reserved slots to transmit data. As the stations are part of a broadcast system, the bit-map is available to all stations to enable them to judge when they should transmit their data in the main part of the operating cycle.

The application of a bit-map protocol to a star network does not, by itself, appear to have any significant advantages. However, the inventor has found that modified bit-map protocols can be used to advantage on star networks. Indeed, it has been found that advantages can still be achieved using protocols similar to bit-map protocols but without an explicit bit-map; furthermore, the protocols under consideration can be used to advantage on network topologies additional to star topologies.

According to the present invention there is provided a communications system having a plurality of stations interconnected by transmission paths for the broadcast transmission of data therebetween, the system being arranged to operate cyclically with each station being permitted to transmit data each cycle and the order of permitted transmission among the stations in each cycle being predetermined; each station when transmitting its data including information indicating the identity of the station, and any one said station with data to transmit in a said cycle being operative:

a) to determine whether it is the first station in the sequence of stations with data to transmit in that cycle and, if so, thereafter to transmit its data, and

b) to determine, where it is not said first station, when it is the next station with data to transmit, this determination involving monitoring received data transmissions to identify the originating station, said one station upon determining that it is the next station to transmit, transmitting its data at a time dependent on the end of transmission by the immediately preceding station in said sequence of stations;

characterized in that where said one station determines that it is the next station to transmit in a cycle, then at least in the case where said immediately-preceding station in said sequence is the station immediately preceding said one station in said order of permitted transmission, said one station is operative to begin transmitting data at a time before it has finished receiving data from said immediately-preceding station, which is judged by said one station will avoid its transmissions colliding at any other station with the transmissions from said immediately-preceding station.

With such an arrangement, two data transmissions may exist on the network at a given time which permits an increase in efficiency over protocols where only one transmission can validly exist at a time. Certain existing protocols for ring or dual-bus networks provide for the coexistence of two or more transmissions on the network at the same time; however such protocols rely on the

inherent characteristics of these topologies and cannot be readily transferred for use with other topologies such as star networks. In contrast, the present invention is suitable for use with star networks.

It is in fact known to have two transmission present at the same time on different branches of a tree-configured telephone network. In the known arrangement, the local telephone exchange determines the distance from the exchange of each user and then asks each user to transmit its data at such a time as to avoid collisions with the other users but possibly at a time to result in more than one transmission being simultaneously present on the network. It will be recognized that this system which is based upon a central intelligence, is wholly different from that of the present invention where each station has its own intelligence and works as its own master rather than to the commands of a central controller.

In the system of the present invention, the said one station in judging the time to initiate its data transmission, is preferably operative:

- a) to ascertain how far off is the end of data reception from said immediately-preceding station in said sequence; and
- b) to initiate transmission of its data at a time preceding said end of data reception by an amount, no greater than the time:

$$t_1 + t_0 + t_2$$

where:

- t_1 is the time taken for a data transmission from said one station to pass along an outward said transmission path from the station to reach any said transmission path, hereinafter the common path, used by said immediately-preceding station in said sequence in transmitting its data to all other stations except said one station,
- t_2 is the time taken for a data transmission from said immediately-preceding station to pass from said common path to said one station along an inward said transmission path thereto, and
- t_0 is the time taken for a data transmission from said immediately-preceding station to pass along the common path from said outward path to said inward path, the time t_0 being negative where the data transmission from said immediately-preceding station passes said outward path before said inward path.

By using fixed-size data packets to transmit data, or variable size data packets with a header containing size information, it becomes a relatively simple task for a station to ascertain when the reception of a current transmission will end.

In the case of a star network with a central

node, collision of the transmissions from one station with those from the immediately-preceding station in the transmission sequence, can be avoided if said one station defers start of transmission to a time preceding the end of data reception from said immediately preceding station by an amount no greater than the round trip time for a transmission from the station to return via the central node.

Preferably, to facilitate the determination by said one station of whether it is the first station in said sequence or when it is the next-to-transmit station, the system is arranged to operate a bit-mapped protocol. More particularly, the system is preferably such that each cycle has an initial reservation phase during which each station is arranged to indicate in a corresponding reservation time slot if it is to transmit data in that cycle, said one station being operative to monitor these reservation indications whereby to determine whether it is said first station in said sequence and, if not, to determine the identity of said immediately-preceding station in said sequence, said one station being further operative when not said first station to determine when it is the station next to transmit by identifying the receipt of data from said immediately-preceding station.

A similar scheme may also be applied to the acknowledgement of successful data transmission. The purpose of transmitting data over the network is, of course, to send it to a particular station on the network. When the data is received by that other station, the receiving station should have some method of acknowledging successful receipt. Therefore, each cycle may further have an acknowledgement phase appearing in the cycle after the data transmissions. When each station has successfully received data from another station, it can acknowledge this by setting a reservation time slot in the reservation phase, this time slot being one associated with that other station, whereby to indicate to the network that the data has been transmitted successfully.

The use of reservation slots may be further expanded. In the application of the present invention to some networks, it may be found that some stations need as low an access delay as possible, and this may be compromised by the fact that a station which has not made a reservation cannot transmit data until the next cycle. Provision can therefore be made for at least some of the stations automatically to make a reservation irrespective of whether they have data to transmit or not. Then, if a station which has automatically made a reservation does not have data to transmit, this must be recognized by the next successive station which does have data to transmit. This may then lead to delays within the cycle, but since automatic reservation is of primary benefit to those stations

which have much data to transmit, it is likely that the delays will not prove a problem in practice.

The determination by each station of when it is to transmit within the sequence of stations with data to transmit, does not necessarily require the use of a bit-mapped protocol with a reservation scheme. As an alternative, it is possible for each station to detect if the immediately preceding station, and indeed successive preceding stations before that, are to transmit, on the basis of detection of the presence or absence of their data transmissions. Where the station detects that the immediately preceding station, and successively preceding stations, have not transmitted data, it may then promptly on that recognition, transmit its own data. Of course, the delays between the data transmission may then be longer, degrading the efficiency of utilisation. However, access delays may be reduced and the implementation is simpler. Therefore, in some situations, this alternative may be preferred.

An embodiment of the present invention will now be described in detail, by way of example, with reference to the accompanying drawings in which;

Figure 1 shows a star network;

Figure 2 shows one cycle of operation for a network operating according to the present invention;

Figure 3 shows a second alternative cycle for a network operating according to the present invention;

Figure 4 shows a third alternative cycle for a network according to the present invention; and Figure 5 shows a star network with couplers linking groups of stations to the central star node;

Figure 6A shows the data paths of the central node 101 of the networks of Figures 1 and 5;

Figure 6B shows the data paths of the couplers of the network of Figure 5; and

Figure 7A and 7B are diagrams of the signal paths near a station in a general network.

Referring first to Figure 1 a star network comprises a plurality of stations 100 connected to a central node 101. Each of those stations 100 may be, for example, a digital computer.

The network shown in Figure 1 may be, for example, a passive optical star network with a passive star node 101 of the general form shown in Figure 6A. The node 101 of Figure 6A has a plurality of input lines 200 coming from the stations 100 and a plurality of output lines 201 leading to the stations 100. A signal on any one of the input lines 200 will be distributed to all output lines 201.

Such passive optical star networks have important advantages over active star, ring, or bus. For example, although the passive star has a potential

single point of failure (at the central passive star coupler node 101) it requires no power, is immune to electromagnetic interference, and has no mechanical or electrical components. It is therefore extremely reliable and does not need optical transceivers, switches or cables. Furthermore, by reducing the number of optical components, passive star optical networks may be constructed more cheaply than other optical networks. Additionally, many localised applications suit a star topology better than a bus or ring. For example, it is frequently the case that each user of a network, or work group, is connected to a central wiring point for e.g. power. The present invention is therefore particularly applicable to passive optical star networks, although it is also applicable to passive electrical star networks, active star networks and other topologies.

In the present invention, as applied to the star network shown in Figure 1, the network operates in a succession of cycles in each of which each station 100 is permitted to transmit one data packet through the network with the order of permitted transmission among the stations in each cycle being predetermined. The stations 100 time their transmission such as to avoid packet collisions at the receiving stations. Since the primary potential point of collision of such data packets is central node 101, it is the time of arrival of the packets at that central node 101 which is best examined when considering how to render transmission over the network collision free.

Of course, one possible scheme is for the station next-to-transmit in the sequence of stations with data to transmit in a cycle, to begin its transmission only after it has heard the immediately preceding station in said sequence complete its transmission. However, since there is inevitably a delay between transmission of a data packet by a station 100 and its reception at the central node 101, waiting until the end of transmission by the immediately-preceding station, is inefficient (it should be noted that the distances between the station 100 and the central node 101 may be more than a kilometre, and so those delay times could become significant). Accordingly, each of the stations of the present network is so arranged that when it has a data packet to transmit, it will start transmitting prior to the end of receipt of the packet sent by the station immediately preceding it in said sequence by an amount judged to just avoid collision at the central node 101. In the simple network of Figure 1, this amount is just less than the round trip time between the station 100 and central node 101.

To implement this scheme, each station is arranged measure how far it is (in terms of transmission time) from the central node 101.

The round trip delay time is dependent on the electronics of the system, the fibre optic transmitter, the fibre itself, and any intermediate optical couplers (not shown). Whilst it may be possible accurately to control the length of the fibre, the delay through the other components cannot be known sufficiently precisely for measurement to an accuracy of one bit. However, it is feasible for a station to measure within a certain tolerance how many bit times it is from the centre.

The format of one cycle which makes use of the present invention is shown in Figure 2. The Figure 2 cycle is made up of two phases, namely a reservation phase 110 and a data phase 111. The reservation phase 110 is notionally divided into a plurality of reservation time slots 112, the number of reservation slots corresponding to the number of stations. As already noted, the stations operate in a predetermined order, there is thus a "first" station of the cycle, a "second" station of the cycle etc., up to an nth station. There are thus a corresponding number of reservation slots 112, and each station may generate an indication in its corresponding reservation slot to indicate that it is to transmit a packet in that cycle. This indication of a packet to transmit is depicted by a logic "1" in Figure 2, with a contrary indication being depicted by logic "0". A "1" indication in a reservation slot effectively reserves a slot for the corresponding station to transmit in the following data phase of the cycle.

In theory, such reservation slots could be only one bit wide, if the station could determine exactly when a signal transmitted was received at the central node 101 of the network. In practice this is not possible, and the notional width of each reservation slot is therefore of a predetermined window length, which window length is determined by the uncertainties in the transmission. In practice, such measurement accuracy may be accurate to around 5 to 6 bits, and thus the window length may be approximately 10 bits, the bit length being determined by a carrier of e.g. square waves transmitted around the system. In order to provide a "1" indication (packet to transmit) in a reservation slot, the corresponding station may apply a pulse to that square wave for more than one bit, for example 2 to 3 bits. This corrupts the square wave signal and this can be detected by standard methods by the other stations.

After the reservation phase 110, the cycle shown in Figure 2 proceeds with its data phase 111. In this phase 111, each station which has reserved a slot in the reservation phase 110 transmits a packet, the transmission of packets being determined by the ordering of the stations. By knowing the predetermined order of station transmission within a cycle and by monitoring the set-

ting of the reservation slots, each station can determine which station, if any, in the sequence of stations that are to transmit a packet during the current cycle, immediately precedes the station in question. Where a station determines it is the first station in said sequence of stations with data to transmit, it can commence its transmission. Each other station in said sequence must listen to the data packets being transmitted over the network and upon receiving the packet from the immediately-preceding station in said sequence, the station can transmit its data packet such as to be promptly received at the central node after the end of the packet from the preceding station. Of course, this assumes that each data packet carries an indicator of its originating station; however, this is standard practice in packet-based networks. Additionally, the station must have some pre-knowledge of the length of a packet so it can judge when the packet from the immediately preceding station in said sequence will end and then time its transmit start (which precedes the end of the preceding packet) accordingly. Where the packets are of fixed length, it is easy for each station to calculate from the start of the packet when it will end, and time its own transmission of a packet accordingly. However, if variable length packets are used, later stations in the cycle must determine from information in the packet itself when it will end, i.e. the packet itself contains information about its length.

It should be noted that there is a minimum packet length also referred to as minimum slot length F_{min} . This minimum slot length is equal to the maximum propagation delay between the stations and the central node. If a station transmits a packet shorter than the minimum slot length, the time difference may be wasted, because the station next to transmit may not be able to detect the end of that packet sufficiently rapidly.

The minimum slot length F_{min} is determined by the maximum length of the link from any of the stations to the central node L_{max} , the bit rate B , and the propagation speed C along the transmission path (e.g. optical fibre) of the star network. Thus:

$$F_{min} = \frac{2BL_{max}}{C}$$

If packets are of a variable length a header of length H will be required in addition to the minimum slot length, as well as the time L_{read} taken by station to interpret the packet length. The effective minimum slot length is now:

$$F_{min} = \frac{2BL_{max}}{C} + H + L_{read}$$

In practice, there must also be a delay between the end of the reservation phase and the start of the first packet within the data phase, and this must be at least twice L_{max} .

The cycle shown in Figure 2 may further be developed, as shown in Figure 3. In this case, the cycle still contains a reservation phase 110 and a data phase 111, but further contains an acknowledgement phase 113. The acknowledgement section 113 notionally contains acknowledgement slots 114, there being as many acknowledgement slots 114 as there are stations. It also can be seen that the acknowledgement phase 113 is separated from the data phase 111 by a delay S_1 . The delay S_1 must be sufficient to give the recipient of the last packet long enough to reply in the next acknowledgement slot 114; further, if only the first station transmits in a cycle, then S_1 must be greater than F_{min} . In practice, S_1 will need to be greater than the lower limits stated in the last sentence as the recipient must have time to check the validity of the packet.

Similarly, there is a delay S_2 after the end of the acknowledgement phase 113 before the next cycle can start. This delay gives the station long enough to reserve a slot in the next cycle if the packet which is applied to the cycle just ending was not acknowledged. If the ordering of the acknowledgement slots 114 is the same as the ordering of the reservation slots 112, then $s_2 > F_{min} \cdot N W_L$ where W_L is the window length between the data packets and N is the number of packets. To acknowledge receipt of a data packet, the station receiving that data packet determines from the data packet which station applied the packet to the cycle and sets an indication in the acknowledgement slot 114 corresponding to the originating station. The originating station may thus detect that the packet has been successfully transmitted, by checking on the indication in its own acknowledgement slot.

As was mentioned above, reservations may be determined by a square wave signal applied to the network. In the arrangement shown in Figure 3, this can be readily achieved by having the last station to transmit apply a square wave to the cycle after the end of transmission of its data packet. This square wave transmission is maintained for the acknowledgement phase 113 and the reservation phase 110 of the next cycle and is terminated at the start of the transmission of a data packet in the next cycle. As was mentioned above, the setting of the slots is then achieved by the corruption of the square wave signal. In practice, it may be simplest to arrange for the square wave to be generated by the last station in the predetermined order rather than the last station actually to transmit in a cycle.

Where a station does not set a "1" indication in

its reservation slot, it cannot transmit a data packet in that cycle. Since the cycle length is primarily determined by the data phase 111 (the relative length of the reservation phase 110 and the acknowledgement phase 113 being very much shorter), it can be seen that if one of the early stations in the predetermined order realises soon after the reservation phase 110 has passed that it wishes to transmit, it must wait for one whole cycle. However, if a station near the end of the predetermined order realises it has data to transmit at a similar point in the cycle, it must wait for nearly two cycles before it can transmit its packet (i.e. one whole cycle plus that part of a cycle occupied by the data packets of the preceding stations in the next cycle). There is thus a variation in the access delay from one station to another. This may be of particular importance where some stations require low access delays. In this case, the arrangement shown in Figure 3 may further be modified by permitting stations requiring low access states automatically to reserve slots every cycle, regardless of whether those stations will use the slots. In this arrangement, if a station has already reserved a slot, it will be able to transmit a data packet at any time up to its data packet slot, and therefore access delay is decreased. The station would always be able to transmit in the current cycle, rather than wait for the next.

The disadvantage of this is that each station must then be able to detect whether a station which has automatically reserved a slot is actually going to apply a data packet to the cycle, and if not, to advance its own transmission of a data packet accordingly. There may therefore be longer delays between one data packet and the next during the data phase 111.

Indeed, this process may be taken further, as shown in Figure 4. In this arrangement, there are no reservation slots, and each station is permitted to transmit a data packet at a predetermined point in the cycle. If a station recognises that the previous station, and indeed any station preceding that, is not going to transmit a data packet in the cycle, then it may advance the timing of its transmission of a data packet accordingly; in this case, where the middle of three stations has nothing to transmit the gap at the central node 101 between the packets from the other stations will be at least F_{min} . Note that since each station does not know in advance which stations have data to transmit, the only circumstance in which it may start transmitting prior to having finished receiving a packet, is when the station recognises that the packet it is receiving is from the station immediately preceding it in said predetermined order.

The Figure 4 arrangement reduces the maximum access delay, and the implementation is sim-

pler, particularly if the acknowledgement phase 113 was also omitted, as shown in Figure 4. However, there is a reduction in the utilisation efficiency because of the potential increased delay between one data packet and the next. Furthermore, in any practical system, in order to ensure synchronization of all stations to the system cycle, there needs to be some transmissions in each cycle even if no stations have data to send. This could be managed, for example, by requiring each station with no data to send to transmit an empty packet of length F_{min} or less thereby flagging progress of the cycle to all other stations; alternatively, whenever a station recognises that its successor has not transmitted a packet, then it must do so itself even if this packet is empty.

Figure 5 shows a modified form of the Figure 1 network in which couplers 102 have been introduced between the stations 100 and the central node 101 to reduce the data flow paths. It can be seen that between the stations 100 and the central node 101, there are couplers 102 to reduce the data flow paths. Thus, in transmission of a data packet from one station 100 to another station 100, data is transmitted from that station 100 to the immediate adjacent coupler 102, from that coupler 102 to the central node 101, from that central node 101 to the coupler 102 of the destination station, and then to the destination station itself. In order to maintain the topography of the network (that is, a central star node through which all packets are routed), each coupler 102 is constructed differently from the central node 101 (shown in Figure 6A). More particularly, as is illustrated in Figure 6B, each coupler has two parts, 202 and 203. The part 202 receives data packets and other signals from the stations 100 via lines 204 and transmits them via line 205 to the central node 101, whose direction is indicated by the arrow 206. Similarly the part 203 receives data packets and other signals from the central station 111 via line 207 and relays those signals to the stations via lines 28. It can thus be seen that signals coming from the stations are routed to the central node 101 only and do not collide with the signals coming from the central node 101. Note that the timing criteria for avoiding collisions at the central node 101 previously discussed with respect to the Figure 1 network will effectively still apply (thus the maximum amount by which a station can advance transmission to proceed the end of receipt of data from the immediately-preceding transmitting station is still the round trip delay time).

Although it is possible for the above described network to be operated with all stations 100 being started at the start-up of the network, it is also possible for a station 100 to join the network after it is running. Once the station has established the

timing of the cycle, it can reserve a slot (e.g. the last empty slot) in the reservation phase of the cycle, and then use its reserved slot in the data phase to signal to the central node 101 that it has joined the network. The central node 101 can then determine a position for that station in the predetermined order of transmission, and the station can then operate normally. Such an arrangement does, of course, require the central node to be provided with intelligence; alternatively, one of the stations 100 can perform this role.

The above described network arrangement permits stations to regain synchronization quickly following errors on the network, with minimum data loss. There are several different errors possible. Firstly, the header of each packet should contain information defining the length of the packet, the station which originated the packet, and the address of the station which is to receive the packet. Failures in the addresses can be detected readily by known systems. Errors of information defining the length of the packet will cause more problems. A station may interpret this information to mean that a packet from an early station is shorter than it really is. If this occurs, and the station is the next to transmit, it will begin to transmit prematurely, and the packets will collide. There will therefore be no gap between the packets from those stations, and the station which is next to transmit will not recognise that a new packet has started. That subsequent station will interpret the collided packets as being from the first of the stations which originated the collided packets, and therefore will wait for the other station which originated the collided packet to arrive. If, however, the maximum packet length is exceeded, the subsequent station can then automatically begin to transmit. Therefore, only two packets are corrupted and the error does not propagate further in the network. Where a station determines that a packet is longer than it really is, no collision or corruption occurs, and there is only a lengthened delay before transmission.

Errors can also occur in the setting in the reservation slots. If a reservation slot is apparently set when no real reservation has been made, then this will cause a delay but will not corrupt data. If, however, a station makes a reservation which is not detected by one or more other stations, then problems can occur. Consider the situation in which a station A fails to recognise a reservation made by station B. If station A does not wish to transmit in the current cycle, it will merely see a cycle longer than it had expected. No data loss will occur. If station A does wish to transmit in the current cycle, and is arranged to transmit its data packet in the cycle before station B, station A will again see a cycle longer than it expected but no data loss will occur. If station A wishes to transmit in the current

cycle and is arranged to transmit its data packet immediately after station B, then a collision will occur. The situation will then be the same as discussed previously for errors in the header length, with a collision between two data packets occurring, those two data packets being lost, but subsequent stations being unaffected after a delay. The situation is similar if station A is to transmit in the current cycle, and is arranged to transmit its data packet in the cycle some time, but not immediately, after station B. In this case a packet from station A will collide with another packet, but not with a packet from station B, otherwise, this situation is the same as before.

Errors can also occur in the acknowledgement slots. If acknowledgement does not occur the packet is retransmitted and provided the probability of this is arranged to be low by suitable protocols, then the delays imposed on the system are not great. The situation where a packet is unsuccessfully transmitted, but the appropriate acknowledgement bit is set, is more rare because it depends on two errors and for a normal system the probability of this is sufficiently low to be ignored.

Thus, the above described network arrangements provide a protocol for use in a communications system, and in a communications method, which is particularly, but not exclusively, adapted to star networks, and allows those networks to operate with a higher efficiency than using known protocols. Although there is some "unfairness" between the stations, with a station later in the predetermined order having longer access delays, this problem only exists at relatively low packet rates, because when the network approaches overload, stations have packets to transmit in every cycle and there is therefore no unfairness. With the present arrangements, errors do not propagate in the network so the network can continue to operate with little data loss. Effectively, the protocols of the above described network arrangements permit more than one successful packet to be in transit at any one time, which is not true of existing protocols applied to star networks.

Although in the embodiments of the invention described above, each station is only allowed to transmit one packet (of fixed or variable length) each cycle, it is possible to arrange for a station to transmit several packets during its turn (or slot) in the data phase of one cycle provided that an appropriate mechanism is established enabling the station next to transmit in the cycle to make a pre-judgment regarding the end of data transmission by the station. Such a mechanism may either be by establishing a predetermined structure for data from the station so that the next station can be pre-programmed with this information, or by including information on the amount of data being trans-

mitted, or the amount remaining to be transmitted, within the data itself.

Another possible refinement is to arrange for a station to have two or more associated slots for transmitting in the data phase of a cycle thereby enabling it to transmit more than once in the cycle. Such an arrangement is conceptually equivalent to having two (or more) stations at the same station address but each with its own slot in the cycle.

For embodiments of the present invention that utilize non-star topologies, the maximum amount by which a station can advance its transmission relative to the end of receipt of data from the immediately-preceding transmitting station, requires to be expressed in more general terms than the round-trip delay time between a station and a central node. Figure 7 illustrates the main factors to be considered in deriving a maximum advance value. In Figure 7A, stations 100A and 100B are two successively transmitting stations in a cycle (station 100B transmitting after station 100A), path 700 is a common data path taken by packets transmitted by stations 100A, 100B on their way to other stations in the network, path 701 is the outward signal path from station 100B, and path 702 is the inward signal path to station 100B. Time t_0 is the time taken for a signal to travel the length of the common path 700 between its junctions with paths 701 and 702, time t_1 is the time taken for a signal to travel the length of the path 701, and time t_2 is the time taken for a signal to travel the length of the path 702. For a packet transmitted by station 100B not to collide with a preceding packet transmitted by station 100A, the maximum advance that station 100B can apply is $(t_1 + t_0 + t_2)$; at this advance the front of the packet from station 100B will join the common path 700 just after the end of the packet from station 100A has passed the junction of path 701 with the common path 700. However, if station 100A was in fact positioned as shown in Figure 7B, then the maximum advance could only be $(t_1 - t_0 + t_2)$ which could be negative indicating that the station 100B needed to delay sending out its data packet beyond the time it detected the end of the packet from station 100A. Of course, the position of station 100A relative to paths 701 and 702 will frequently not be known to station 100B so that the worst case must be assumed. For certain topologies it is possible to derive simplified expressions for the maximum advance; one such case is, of course, a star topology. Thus for the Figure 1 network, the maximum advance is $(t_1 + t_2)$ as t_0 will be substantially zero. For the Figure 5 network, in the case where stations 100A and 100B are connected to the same coupler 102 and the common path 700 extends from the coupler to the central node and back, the station 100A will in fact always be in the relation-

ship of Figure 7A to station 100B so that the maximum advance is ($t_1 + t_0 + t_2$) which is equivalent to the round trip delay from station 100B to the central node. Other relative positionings of the stations 100A and 100B in the Figure 5 network produce the same result.

Claims

1. A communications system having a plurality of stations (100) interconnected by transmission paths for the broadcast transmission of data therebetween, the system being arranged to operate cyclically with each station (100) being permitted to transmit data each cycle and the order of permitted transmission among the stations in each cycle being predetermined; each station when transmitting its data including information indicating the identity of the station, and any one said station (100) with data to transmit in a said cycle being operative:

a) to determine whether it is the first station in the sequence of stations with data to transmit in that cycle and, if so, thereafter to transmit its data, and

b) to determine, where it is not said first station, when it is the next station with data to transmit, this determination involving monitoring received data transmissions to identify the originating station, said one station upon determining that it is the next station to transmit, transmitting its data at a time dependent on the end of transmission by the immediately preceding station in said sequence of stations;

characterised in that where said one station determines that it is the next station to transmit in a cycle, then at least in the case where said immediately-preceding station in said sequence is the station immediately preceding said one station in said order of permitted transmission, said one station is operative to begin transmitting data at a time before it has finished receiving data from said immediately-preceding station, which is judged by said one station will avoid its transmissions colliding at any other station with the transmissions from said immediately-preceding station.

2. A communications system according to claim 1, wherein said one station in judging the time to initiate its data transmission, is operative:

a) to ascertain how far off is the end of data reception from said immediately-preceding station in said sequence; and

b) to initiate transmission of its data at a time preceding said end of data reception by an amount, no greater than the time:

$t_1 + t_0 + t_2$

where:

- t_1 is the time taken for a data transmission from said one station to pass along an outward said transmission path from the station to reach any said transmission path, hereinafter the common path, used by said immediately-preceding station in said sequence in transmitting its data to all other stations except said one station,
- t_2 is the time taken for a data transmission from said immediately-preceding station to pass from said common path to said one station along an inward said transmission path thereto, and
- t_0 is the time taken for a data transmission from said immediately-preceding station to pass along the common path from said outward path to said inward path, the time t_0 being negative where the data transmission from said immediately-preceding station passes said outward path before said inward path.

3. A communications system according to claim 2, wherein each station is arranged to transmit a fixed amount of data during each cycle that it transmits data, said one station being operative to ascertain said end of data reception based on a determination of the progress of reception of data from said immediately-preceding station and a prior knowledge of the size of said fixed amount of data.

4. A communications system according to claim 2, wherein the amount of data transmitted by a said station in a cycle is variable and the station is operative when transmitting its data to include an indicator of the amount of data being transmitted, said one station being operative to ascertain said end of data reception based on a determination of the progress of reception of data from said immediately-preceding station and said indicator of the amount of data being transmitted.

5. A communication system according to any one of the preceding claims, wherein the plurality of stations (100) are interconnected in a star network via a central node (101), said one station initiating transmission of its data at a time preceding said end of data reception from said immediately-preceding station by an amount no greater than the round trip taken for a transmission from the station to return via

- said central node (101).
6. A communications system according to any one of the preceding claims, wherein each cycle has an initial reservation phase during which each station (100) is arranged to indicate in a corresponding reservation time slot (112) if it is to transmit data in that cycle, said one station being operative to monitor these reservation indications whereby to determine whether it is said first station in said sequence and, if not, to determine the identity of said immediately-preceding station in said sequence, said one station being further operative when not said first station to determine when it is the station next to transmit by identifying the receipt of data from said immediately-preceding station.
7. A communications system according to claim 6, wherein said one station, when not said first station in said sequence, is operative to begin transmitting data before it has finished receiving data from said immediately-preceding station in said sequence independently of whether said immediately-preceding station is also the station immediately preceding said one station in said order of permitted transmission.
8. A communications system according to claim 6 or claim 7, wherein at least one of the stations (100) is arranged always to indicate in the corresponding reservation time slot (112) that the at least one of the stations (100) is to transmit data irrespective of whether the at least one of the stations knows, at the time of said reservation phase, if it is to transmit data in the cycle.
9. A communications system according to any one of claims 6 to 8, wherein each cycle has an acknowledgement phase during which any one of the stations (100) that receives data from a second one of the stations (100) is arranged to indicate that reception in an acknowledgement time (114) corresponding to said second one of the stations (100).
10. A communications system according to anyone of claims 1 to 5, wherein said station in determining when it is the next-to-transmit station, monitors progress through the cycle in said order of predetermined transmission by identifying the or each earlier transmitting station in said cycle by its transmission, and by judging that a next station in said order is not going to transmit in said cycle upon a transmission not being received from that station

during a timed period following the end of data reception from the station immediately preceding said next station in said order.

5

10

15

20

25

30

35

40

45

50

55




FIG 1




FIG 2




FIG 3





FIG 4

FIG 5






FIG 6A




FIG 6B




FIG 7A




FIG 7B



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number

EP 90 10 1759

DOCUMENTS CONSIDERED TO BE RELEVANT			CLASSIFICATION OF THE APPLICATION (Int. Cl. 5)
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	
X	EP-A-0 255 442 (ALCATEL N.V.) * Column 3, line 44 - column 4, line 35; figures 1a,1c *	1,2,6	H 04 L 12/44
A	---	3-5,10	
X	IBM TECHNICAL DISCLOSURE BULLETIN, vol. 32, no. 5B, October 1989, pages 357-361, New York, US; "Adaptive slotted bus Metropolitan area network protocol" * Page 357, line 1 - page 358, line 17; page 359, lines 13-33; page 360, lines 1-13; figures 1,2 *	1	
A	IDEM	2-5,6, 10	

			TECHNICAL FIELDS SEARCHED (Int. CL5)
			H 04 L
<p>The present search report has been drawn up for all claims</p>			
Place of search	Date of completion of the search	Examiner	
THE HAGUE	10-09-1990	DE LA FUENTE DEL AGUA P.	
CATEGORY OF CITED DOCUMENTS		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document	
X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document			